
Subject: Re: [RFC][PATCH 4/7] RSS accounting hooks over the code

Posted by [dev](#) on Mon, 12 Mar 2007 16:23:25 GMT

[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Pavel Emelianov <xemul@sw.ru> writes:

>

>

>>Pages are charged to their first touchers which are
>>determined using pages' mapcount manipulations in
>>rmap calls.

>

>

> NAK pages should be charged to every rss group whose mm_struct they
> are mapped into.

For these you essentially need per-container page->_mapcount counter,
otherwise you can't detect whether rss group still has the page in question being mapped
in its processes' address spaces or not.

1. This was discussed before and considered to be ok by all the resource management involved people.
2. this can be done with a-la page beancounters which are used in OVZ for shared fractions accounting. It's a next step forward.

If you know how to get "pages should be charged to every rss group whose mm_struct they are mapped into" w/o additional pointer in struct page, please throw me an idea.

Thanks,
Kirill

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
